

Probabilistic Feature Selection and Classification Vector Machine

Bingbing Jiang, Chang Li, Huanhuan Chen, *Senior Member, IEEE*, and Xin Yao, *Fellow, IEEE* and Maarten de Rijke

Abstract—Sparse Bayesian learning is one of the state-of-the-art machine learning algorithms, which is able to make stable and reliable probabilistic predictions. However, some of these algorithms, e.g. probabilistic classification vector machine (PCVM) and relevant vector machine (RVM), are not capable of eliminating irrelevant and redundant features which could lead to performance degradation. To tackle this problem, in this paper, we propose a sparse Bayesian classifier which simultaneously selects the relevant samples and features. We name this classifier a probabilistic feature selection and classification vector machine (PFCVM), in which truncated Gaussian distributions are employed as both sample and feature priors. In order to derive the analytical solution for the proposed algorithm, we use Laplace approximation to calculate approximate posteriors and marginal likelihoods. Finally, we obtain the optimized parameters and hyperparameters by the type-II maximum likelihood method. The experiments on synthetic data set, benchmark data sets and high dimensional data sets validate the performance of PFCVM under two criteria: accuracy of classification and efficacy of selected features. Finally, we analyze the generalization performance of PFCVM and derive a generalization error bound for PFCVM. Then by tightening the bound, we demonstrate the significance of the sparseness for the model.

Index Terms—Sparse Bayesian learning, feature selection, probabilistic classification model, Laplace approximation, generalization bound.

I. INTRODUCTION

In supervised learning, we are given input feature vectors $\mathbf{x} = \{x_i \in \mathbb{R}^M\}_{i=1}^N$ and corresponding labels $\mathbf{y} = \{y_i\}_{i=1}^N$.¹ The goal is to predict the label of a new datum \hat{x} based on the training data set $\mathbf{S} = \{\mathbf{x}, \mathbf{y}\}$ together with other prior knowledge. For regression, we are given continuous labels $y \in \mathbb{R}$, while for classification we are given discrete labels. In this paper we focus on the binary classification case, in which $y \in \{-1, +1\}$.

Recently, learning sparseness from large-scale datasets has generated significant research interest [1–5]. Among the methods prooposed, the support vector machine (SVM) [1], which based on the kernel trick [6] to create a non-linear decision

boundary with a small number of support vectors, is the state-of-the-art algorithm. The prediction function of SVM is a combination of basis functions:²

$$f(\hat{x}; \mathbf{w}) = \sum_{i=1}^N \phi(\hat{x}, x_i) w_i + b, \quad (1)$$

where $\phi(\cdot, \cdot)$ is the basis function, $\mathbf{w} = \{w_i\}_{i=1}^N$ are sample weights, and b is the bias.

Similar to SVM, many sparse Bayesian algorithms use (1) as their decision function; examples include like the relevance vector machine [7, RVM] and the probabilistic classification vector machine [8, PCVM]. Unlike SVMs, whose weights are determined by maximizing the decision margin and limited to hard binary classification, Bayesian algorithms optimize the parameters by computing the maximum probability and make predictions based on the average of the prediction function over the posterior. For example, PCVMs compute the maximum a posteriori (MAP) estimation using the expectation-maximization (EM) algorithm; RVMs and efficient probabilistic classification vector machines [3, EPCVM] compute the type-II maximum likelihood [9] to obtain the best parameters. These algorithms have to deal with different scales of features.

In order to fit different scales of features, basis functions are always controlled by basis parameters (or kernel parameters). For example, in LIBSVM [10] with Gaussian radial basis functions (RBF) $\phi(x, z) = \exp(-\vartheta \|x - z\|^2)$, the default ϑ is set relatively small for high-dimensional datasets, and large for low-dimensional datasets. Although the use of basis parameters may help to address the curse of dimensionality [11], the performance might be degraded when there are lots of irrelevant and/or redundant features [2, 12, 13]. Parameterized basis functions are designed to deal with this problem. There are two popular basis functions that can incorporate feature parameters easily:

Gaussian RBF

$$\phi_{\theta}(x, z) = \exp\left(-\sum_{k=1}^M \theta_k (x^k - z^k)^2\right), \quad (2)$$

Pth order polynomial:

$$\phi_{\theta}(x, z) = (1 + \sum_{k=1}^M \theta_k x^k z^k)^P, \quad (3)$$

²In the rest of this paper, we prefer to use the term of basis function instead of kernel function, because, except for SVM, the basis functions used in this paper are free of Mercer's condition.

Bingbing Jiang, Huanhuan Chen, and Xin Yao are with UBRI, School of Computer Science and Technology, University of Science and Technology of China (USTC), Hefei, 230027, China, email: jiangbb@mail.ustc.edu.cn, hchen@ustc.edu.cn. Xin Yao is also with CERCIA, School of Computer Science, University of Birmingham, Birmingham B15 2TT, United Kingdom, email: X.Yao@cs.bham.ac.uk. Chang Li and Maarten de Rijke are with the Informatics Institute, University of Amsterdam, 1098 XH Amsterdam, The Netherlands, email: {c.li, derijke}@uva.nl

¹In this paper, the subscript of a sample x , i.e. x_i , denotes the i th sample and the superscript of a sample x , i.e. x^k , denotes the k th dimension.

where the subscript denotes the corresponding index of features, and $\theta \in \mathbb{R}^M$ are feature parameters (also called feature weights). Once a feature weight $\theta_k \rightarrow 0$,³ the corresponding feature will not contribute to the classification.

Feature selection plays an important role in reducing the dimensionality of data. Feature selection methods can be divided into three groups: *filter methods* [14–16], *wrapper methods* [17], and *embedded methods* [2, 5, 12, 13, 18]. Filter methods independently select the subset of features from the classifier learning. Wrapper methods first train the classifier with all possible feature subsets and then select a specific subset with the best performance. Therefore, filter methods have relatively low accuracy but high computational efficiency compared to wrapper methods. Embedded methods embed feature selection in the training process, which aims to combine the advantages of the filter and wrapper methods. Peng et al. [14], propose a filter method, mRMR, that selects relevant features and simultaneously removes redundant features according to mutual information. Based on the above basis functions, Weston et al. [17] develop a wrapper feature selection schema for SVM, and Nguyen and De la Torre [13] design an embedded feature selection model, Weight SVM (WSVM), that can jointly perform feature selection and classifier learning for linear and non-linear SVM. However, these methods are not able to produce probabilistic outputs, and they need to predetermine the number of selected features before training.

For Bayesian (embedded) approaches, a joint classifier and feature optimization algorithm (JCFO⁴) is proposed in [12]; the authors adopt a sparse Bayesian model to simultaneously perform classifier learning and feature selection. To select relevant features, JCFO introduces hierarchical sparseness promoting priors on feature weights and then employs EM and gradient based methods to optimize the feature weights. In order to simultaneously select relevance samples and features, Mohsenzadeh et al. [18] extend the standard RVM and then design the relevance sample feature machine (RSFM) and an incrementally learning version (IRSFM) [5], that scales the basis parameters in RVM to a vector and applies zero-mean Gaussian priors on feature weights to generate sparsity in the feature space. Li and Chen [2] propose an EM algorithm based joint feature selection strategy for PCVM (denoted as PFCVM_{EM}), in which they add truncated Gaussian priors to features to enable PCVM to jointly select relevant samples and features. However, JCFO, PFCVM_{EM} and RSFM use an EM algorithm to calculate a maximum a posteriori point estimate of the sample and feature parameters. As pointed out by Chen et al. [3], the EM algorithm has the following limitations: first, it is sensitive to the starting points and cannot guarantee convergence to global maxima or minima; second, the EM algorithm results in a MAP point estimate, which limits to the Bayes estimator with the 0-1 loss function and cannot represent all advantages of the Bayesian framework.

JCFO, RSFM and IRSFM adopt a zero-mean Gaussian prior distribution over sample weights, and RSFM and IRSFM also

use this prior distribution over feature weights. As a result of adopting a zero-mean Gaussian prior over samples, some training samples that belong to the positive class ($y_i = +1$) will receive negative weights and vice versa; this may result in instability and degeneration in solutions [8]. Also, for RSFM and IRSFM, zero-mean Gaussian feature priors will lead to negative feature weights, which reduces the value of kernel functions for two samples the similarity in the correspond features is increased [12]. Finally, RSFM and IRSFM a $N \times M$ kernel matrix for each sample, which yields a space complexity of at least $O(N^2M)$ to store the design kernel matrices.

We propose a complete, sparse Bayesian method, i.e., a Laplace approximation based feature selection method PCVM (PFCVM_{LP}), which uses the type-II maximum likelihood method to arrive at a fully Bayesian estimation. In contrast to filter methods such as mRMR[14] and the student t-test [15], and the embedded WSVM method [13], the proposed PFCVM_{LP} method can make probabilistic predictions with adaptive sparsity in the sample and feature space. Moreover, PFCVM_{LP} adopts truncated Gaussian priors as both sample and feature priors, which avoids the negative values for sample and feature weights. We summarize the main contributions as follows.

- Unlike PCVM and RVM, our proposed algorithm can automatically select informative features in non-linear basis functions, which leads to a more robust classifier for high-dimensional data sets.
- Compared with PFCVM_{EM} [2], JCFO [12] and RSFM [18], PFCVM_{LP} computes the type-II maximum likelihood [7], thereby avoiding limitations of the EM algorithm.
- We derive a generalization bound for PFCVM_{LP}. By analyzing the bound, we demonstrate the significance of feature selection and introduce a way of choosing the initial values.

The rest of the paper is structured as follows. Background knowledge on sparse Bayesian learning is introduced in Section II. Section III details the implementation of simultaneously optimizing sample and feature weights of PFCVM_{LP}. In Section IV experiments are designed to evaluate both the accuracy of classification and the effectiveness of feature selection. Analyses of sparsity and generalization for PFCVM_{LP} are presented in Section V. We conclude in Section VI.

II. SPARSE BAYESIAN FRAMEWORK

In the sparse Bayesian learning framework, we usually use the Laplace distribution and/or the student's-t distribution as the sparseness-promoting prior. In binary classification problems, we choose a Bernoulli distribution as the likelihood function. Together with the proper marginal likelihood, we can compute the parameters' distribution (posteriori distribution) either by MAP point estimation or by type-II-maximum likelihood. Below, we detail the implementation of this framework.

³Practically lots of feature weights $\theta_k \rightarrow 0$. When a certain θ_k is smaller than a threshold, we will set it 0.

⁴http://www.cns.atr.jp/~oyamashi/SLR_WEB/

A. Model Specification

We concentrate on a linear combination of basis functions. To simplify our notation, the decision function is defined as:

$$f(\mathbf{x}; \mathbf{w}, \boldsymbol{\theta}) = \Phi_{\boldsymbol{\theta}}(\mathbf{x})\mathbf{w}, \quad (4)$$

where \mathbf{w} denotes the $N + 1$ -dimensional sample weights; w_0 denotes the bias; $\Phi_{\boldsymbol{\theta}}(\mathbf{x})$ is an $N \times (N + 1)$ basis function matrix, except for the first column $\phi_{\boldsymbol{\theta},0}(\mathbf{x}) = [1, \dots, 1]^T$, while other component $\phi_{\boldsymbol{\theta},ij} = \phi_{\boldsymbol{\theta}}(x_i, x_j) \times y_j$,⁵ and $\boldsymbol{\theta} \in \mathbb{R}^M$ is the feature weights.

As probabilistic outputs are continuous values in $[0, 1]$, we need a link function to obtain a smooth transformation from $[-\infty, +\infty]$ to $[0, 1]$. Here, we use a sigmoid function $\sigma(z) = \frac{1}{1+e^{-z}}$ to map (4) to $[0, 1]$. Then, we combine this mapping with a Bernoulli distribution to compute the following likelihood function:

$$p(\mathbf{t} | \mathbf{w}, \boldsymbol{\theta}, \mathbf{S}) = \prod_{i=1}^N \sigma_i^{t_i} (1 - \sigma_i)^{(1-t_i)},$$

where $t_i = (y_i + 1)/2$ denotes the probabilistic target of the i -th sample and σ_i denotes the sigmoid mapping for the i -th sample: $\sigma_i = \sigma(f(x_i; \mathbf{w}, \boldsymbol{\theta}))$. The vector $\mathbf{t} = (t_1, \dots, t_N)^T$ consists of the probabilistic targets of all training samples and $\mathbf{S} = \{\mathbf{x}, \mathbf{y}\}$ is the training set.

B. Priors over Weights and Features

According to Chen et al. [8], a truncated Gaussian prior may result in the proper sparseness to sample weights. Following this idea, we introduce a non-negative left-truncated Gaussian prior $N_t(w_i | 0, \alpha_i^{-1})$ to each sample weight w_i :

$$\begin{aligned} p(w_i | \alpha_i) &= \begin{cases} 2N(w_i | 0, \alpha_i^{-1}) & \text{if } w_i \geq 0 \\ 0 & \text{otherwise} \end{cases} \\ &= 2N(w_i | 0, \alpha_i^{-1}) \cdot 1_{w_i \geq 0}(w_i), \end{aligned} \quad (5)$$

where α_i (precision) is a hyperparameter, which is equal to the inverse of variance, and $1_{x \geq 0}(x)$ is an indicator function that returns 1 for each $x \geq 0$ and 0 otherwise. For the bias w_0 , we introduce a zero-mean Gaussian prior $N(w_0 | 0, \alpha_0^{-1})$:

$$p(w_0 | \alpha_0) = N(w_0 | 0, \alpha_0^{-1}). \quad (6)$$

Assuming that the sample weights are independent and identically distributed (i.i.d.), we can compute the priors over sample weights as follows:

$$\begin{aligned} p(\mathbf{w} | \boldsymbol{\alpha}) &= \prod_{i=0}^N p(w_i | \alpha_i) = N(w_0 | 0, \alpha_0^{-1}) \prod_{i=1}^N N_t(w_i | 0, \alpha_i^{-1}), \end{aligned} \quad (7)$$

where $\boldsymbol{\alpha} = (\alpha_0, \dots, \alpha_N)^T$ and $N_t(w_i | 0, \alpha_i^{-1})$ denotes the left truncated Gaussian distribution.

Feature weights indicate the importance of features. For important features, the corresponding weights are set to relatively large values, and vice versa. For irrelevant and/or redundant

features, the weights are set to 0. Following [12], we should not allow negative values for feature weights. Based on these discussions, we introduce left truncated Gaussian priors for feature weights. Under the i.i.d. assumption, the prior over features is computed as follows:

$$p(\boldsymbol{\theta} | \boldsymbol{\beta}) = \prod_{k=1}^M p(\theta_k | \beta_k) = \prod_{k=1}^M N_t(\theta_k | 0, \beta_k^{-1}),$$

where $\boldsymbol{\beta} = (\beta_1, \dots, \beta_M)^T$ are hyperparameters of feature weights. Each prior is formalized as follows:

$$\begin{aligned} p(\theta_k | \beta_k) &= \begin{cases} 2N(\theta_k | 0, \beta_k^{-1}) & \text{if } \theta_k \geq 0, \\ 0 & \text{otherwise,} \end{cases} \\ &= 2N(\theta_k | 0, \beta_k^{-1}) \cdot 1_{\theta_k \geq 0}(\theta_k). \end{aligned} \quad (8)$$

For both kinds of priors, we introduce Gamma distributions for α_i and β_k as hyperpriors. The truncated Gaussian priors will work together with the flat Gamma hyperpriors and result in truncated hierarchical Student's-t priors over weights. These hierarchical priors, which are similar to Laplace priors, work as L1 regularization and lead to sparse solutions [8, 12].

C. Computing Posteriors

The posterior in a Bayesian framework contains the distribution of all parameters. Computing parameters boils down to updating posteriors. Having priors and likelihood, posteriors can be computed with the following formula:

$$p(\mathbf{w}, \boldsymbol{\theta} | \mathbf{t}, \boldsymbol{\alpha}, \boldsymbol{\beta}) = \frac{p(\mathbf{t} | \mathbf{w}, \boldsymbol{\theta}, \mathbf{S}) p(\mathbf{w} | \boldsymbol{\alpha}) p(\boldsymbol{\theta} | \boldsymbol{\beta})}{p(\mathbf{t} | \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{S})}. \quad (9)$$

Some methods, such as PCVM, PFCVM_{EM} and JCFO, discard the marginal likelihood and use the EM algorithm to obtain an MAP point estimation of parameters. Although an efficient estimation can be obtained by the EM algorithm, it overlooks the information in the marginal likelihood and is not regarded as a complete Bayesian estimation. Other methods, such as RVM and EPCVM, retain the marginal likelihood. They compute the type-II maximum likelihood and obtain a complete Bayesian solution.

The predicted distribution for the new datum \hat{x} is computed as follows:

$$p(\hat{y} | \hat{x}, \mathbf{t}, \boldsymbol{\alpha}, \boldsymbol{\beta}) = \int p(\hat{y} | \hat{x}, \mathbf{w}, \boldsymbol{\theta}) p(\mathbf{w}, \boldsymbol{\theta} | \mathbf{t}, \boldsymbol{\alpha}, \boldsymbol{\beta}) d\mathbf{w} d\boldsymbol{\theta}.$$

If both terms in the integral are Gaussian distributions, it is easy to compute this integral analytically. We will detail the implementation of PFCVM_{LP} in the next section.

III. PROBABILISTIC FEATURE SELECTION CLASSIFICATION VECTOR MACHINE

Details of computing sample weights and sample hyperparameters were reported by Chen et al. [3]. In this section we mainly focus on computing parameters and hyperparameters for features.

⁵We assume that each sample weight has the same sign as the corresponding label. So by multiplying the basis vector with the corresponding label, we can assume that all sample weights are non-negative.

A. Approximations for Posterior Distributions

Since the indicator function in (8) is not differentiable, an approximate function is required to smoothly approximate the indicator function. Here, we use a parameterized sigmoid assumption. Fig. 1 shows the approximation of an indicator function made by a sigmoid function $\sigma(\lambda x)$. As depicted in Fig. 1, the larger λ is, the more accurate approximation a sigmoid function will make. In PFCVM_{LP} , we choose $\sigma(5x)$ as the approximation function.




Fig. 1. Illustration of the indicator function and the sigmoid function.

We calculate (9) by the Laplace approximation, in which the Gaussian distributions⁶ $N(\mathbf{u}_\theta, \Sigma_\theta)$ and $N(\mathbf{u}_w, \Sigma_w)$, are used to approximate the unknown posteriors of feature and sample weights, respectively. We start with the logarithm of Equation (9) by the follow formula:

$$\begin{aligned} Q(\mathbf{w}, \theta) &= \log\{p(\mathbf{t} | \mathbf{w}, \theta, \mathbf{S})p(\mathbf{w} | \alpha)p(\theta | \beta)\} - \log p(\mathbf{t} | \alpha, \beta, \mathbf{S}) \\ &= \sum_{n=1}^N [t_n \log \sigma_n + (1 - t_n) \log(1 - \sigma_n)] - \frac{1}{2} \mathbf{w}^T \mathbf{A} \mathbf{w} \\ &\quad - \frac{1}{2} \theta^T \mathbf{B} \theta + \sum_{i=1}^N \log 1_{w_i \geq 0}(w_i) + \sum_{k=1}^M \log 1_{\theta_k \geq 0}(\theta_k) \\ &\quad + \text{const}, \end{aligned}$$

where $\mathbf{A} = \text{diag}(\alpha_0, \dots, \alpha_N)$, $\mathbf{B} = \text{diag}(\beta_1, \dots, \beta_M)$ and const is independent of \mathbf{w} and θ .

Using the sigmoid approximation, we substitute $1_{x \geq 0}(x)$ by $\sigma(\lambda x)$ with $\lambda = 5$. We can compute the derivative of the feature posterior function as follows:

$$\frac{\partial Q(\mathbf{w}, \theta)}{\partial \theta} = -\mathbf{B}\theta + \mathbf{D}^T(\mathbf{t} - \sigma) + \mathbf{k}_\theta,$$

where $\mathbf{k}_\theta = [\lambda(1 - \sigma(\lambda\theta_1)), \dots, \lambda(1 - \sigma(\lambda\theta_M))]^T$ is an M -dimensional vector, $\sigma = [\sigma_1, \dots, \sigma_N]^T$, and $\mathbf{D} = \frac{\partial \Phi_\theta \mathbf{w}}{\partial \theta}$.

⁶Because of the truncated prior assumption, we should take the positive quadrant part of the two Gaussian distributions, which only have an extra normalization term. Fortunately, the normalization term is independent of \mathbf{w} and θ . So, in the derivation, we still use the Gaussian distributions.

For Gaussian RBF

$$D_{i,k} = - \sum_{j=1}^N w_j \phi_{\theta,ij} (x_i^k - x_j^k)^2.$$

For P th order polynomial

$$D_{i,k} = P x_i^k \sum_{j=1}^N w_j \phi_{\theta,ij}^{(P-1)/P} x_j^k.$$

The mean \mathbf{u}_θ of the feature posterior distribution is calculated by setting $\frac{\partial Q(\mathbf{w}, \theta)}{\partial \theta} = 0$:

$$\mathbf{u}_\theta = \mathbf{B}^{-1} (\mathbf{D}^T(\mathbf{t} - \sigma) + \mathbf{k}_\theta). \quad (10)$$

Then we compute the second-order derivative of $Q(\mathbf{w}, \theta)$, the Hessian matrix:

$$\frac{\partial^2 Q(\mathbf{w}, \theta)}{\partial \theta^2} = -\mathbf{O}_\theta - \mathbf{B} - \mathbf{D}^T \mathbf{C} \mathbf{D} + \mathbf{E},$$

where $\mathbf{O}_\theta = \text{diag}(\lambda^2 \sigma(\lambda\theta_1)(1 - \sigma(\lambda\theta_1)), \dots, \lambda^2 \sigma(\lambda\theta_M)(1 - \sigma(\lambda\theta_M)))$ is an $M \times M$ diagonal matrix, and \mathbf{C} is an $N \times N$ diagonal matrix $\mathbf{C} = \text{diag}((1 - \sigma_1)\sigma_1, \dots, (1 - \sigma_N)\sigma_N)$. \mathbf{E} denotes $\frac{\partial \mathbf{D}^T}{\partial \theta}(\mathbf{t} - \sigma)$ and is computed as follows:

For Gaussian RBF:

$$\begin{aligned} E_{i,k} &= \sum_{p=1}^N \left[(t_p - \sigma_p) \sum_{j=1}^N \phi_{\theta,pj} w_j (x_p^i - x_j^i)^2 \right. \\ &\quad \left. \times (x_p^k - x_j^k)^2 \right]. \end{aligned}$$

For P th order polynomial:

$$\begin{aligned} E_{i,k} &= \sum_{p=1}^N \left[(t_p - \sigma_p) x_p^i x_p^k \sum_{j=1}^N \phi_{\theta,pj}^{(P-2)/P} w_j x_j^i x_j^k \right] \\ &\quad \times P(P-1). \end{aligned}$$

The covariance of this approximate posterior distribution equals the negative inverse of the Hessian matrix:

$$\Sigma_\theta = (\mathbf{D}^T \mathbf{C} \mathbf{D} + \mathbf{B} + \mathbf{O}_\theta - \mathbf{E})^{-1}. \quad (11)$$

Practically, we use Cholesky decomposition to compute the robust inversion.

In the same way, we can obtain \mathbf{u}_w and Σ_w by computing the derivative of $Q(\mathbf{w}, \theta)$ with respect to \mathbf{w} :

$$\mathbf{u}_w = \mathbf{A}^{-1} (\Phi_\theta^T (\mathbf{t} - \sigma) + \mathbf{k}_w) \quad (12)$$

$$\Sigma_w = (\Phi_\theta^T \mathbf{C} \Phi_\theta + \mathbf{A} + \mathbf{O}_w)^{-1}, \quad (13)$$

where $\mathbf{k}_w = [0, \lambda(1 - \sigma(\lambda w_1)), \dots, \beta(1 - \sigma(\lambda w_N))]^T$ is an $(N+1)$ -dimension vector, and $\mathbf{O}_w = \text{diag}(0, \lambda^2 \sigma(\lambda w_1)(1 - \sigma(\lambda w_1)), \dots, \lambda^2 \sigma(\lambda w_N)(1 - \sigma(\lambda w_N)))$ is an $(N+1) \times (N+1)$ diagonal matrix.

After the derivation, the indicator functions degenerate into vectors and matrices, \mathbf{k}_θ in (10), \mathbf{O}_θ in (11) for the feature posterior, and \mathbf{k}_w in (12), and \mathbf{O}_w in (13) for the sample posterior. These two matrices will hold the non-negative

property of the sample and feature weights, which is consistent with the prior assumption.

With the approximated posterior distributions, $N(\mathbf{u}_\theta, \Sigma_\theta)$ and $N(\mathbf{u}_w, \Sigma_w)$, optimizing PFCVM_{LP} boils down to maximizing the posteriori mode of the hyperparameters, which means maximizing $p(\alpha, \beta | \mathbf{t}) \propto p(\mathbf{t} | \alpha, \beta, \mathbf{S})p(\alpha)p(\beta)$ with respect to α and β . As we use flat Gamma distributions over α and β , the maximization depends on the marginal likelihood $p(\mathbf{t} | \alpha, \beta, \mathbf{S})$ [7]. In the next section, the optimal marginal likelihood is obtained through the type-II maximum likelihood method.

B. Maximum Marginal Likelihood

In Bayesian models, the marginal likelihood function is computed as follows:

$$p(\mathbf{t} | \alpha, \beta, \mathbf{S}) = \int p(\mathbf{t} | \mathbf{w}, \theta, \mathbf{S})p(\mathbf{w} | \alpha)p(\theta | \beta)d\mathbf{w}d\theta. \quad (14)$$

However, when the likelihood function is a Bernoulli distribution and the priors are approximated by Gaussian distributions, the maximization of (14) cannot be derived in closed form. Thus we introduce an iterative estimation solution. The details of the hyperparameter optimization and the derivation of maximizing marginal likelihood, are specified in Appendix A. Here, we use the methodology of Bayesian Occam's razor [19]. The update formula of the feature hyperparameters is rearranged and simplified as:

$$\beta_k^{new} = \frac{\gamma_k}{u_{\theta,k}^2}, \quad (15)$$

where $u_{\theta,k}$ is the k -th mean of feature weights in (10), and we denote $\gamma_k \equiv 1 - \beta_k \Sigma_{kk}$, where Σ_{kk} is the k -th diagonal covariance element in (11) and $\beta_k \Sigma_{kk}$ works as Occam's factor, which can automatically find a balanced solution between complexity and accuracy of PFCVM_{LP}. The details of updating the sample hyperparameters α_i are same as for β_k , and we omit them.

In the training step, we will eliminate a feature when the corresponding β_k is larger than a specified threshold. In this case, the feature weight θ_k is dominated by the prior distribution and restricted to a small neighborhood around 0. Hence this feature contributes little to the classification performance. At the start of the iterative process, all samples and features are included in the model. As iterations proceed, N and M are quickly reduced, which accelerates the speed of the iterations. Further analysis of complexity are reported in Section IV-D. In the next subsection, we demonstrate how to make predictions on new data.

C. Making Predictions

When predicting the label of new sample \hat{x} , instead of making a hard binary decision, we prefer to estimate the uncertainty in the decision, the posterior probability of the prediction $p(\hat{y} = 1 | \hat{x}, \mathbf{S})$. Incorporating the Bernoulli likelihood, the Bayesian model enables the sigmoid function $\sigma(f(\hat{x}))$ to be regarded as a consistent estimate of $p(\hat{y} = 1 | \hat{x}, \mathbf{S})$ [7].

We can compute the probability of prediction in the following way:

$$p(\hat{y} = 1 | \hat{x}, \mathbf{S}) = \int p(\hat{y} = 1 | \mathbf{w}, \hat{x}, \mathbf{S})q(\mathbf{w})d\mathbf{w},$$

where $p(\hat{y} = 1 | \mathbf{w}, \hat{x}, \mathbf{S}) = \sigma(\mathbf{u}_w^T \phi_\theta(\hat{x}))$ and $q(\mathbf{w})$ denotes the posterior of sample weights. Employing the posteriori approximation in Section III-A, we have $q(\mathbf{w}) \approx N(\mathbf{w} | \mathbf{u}_w, \Sigma_w)$. According to [20], we have:

$$p(\hat{y} = 1 | \hat{x}, \mathbf{S}) = \int \sigma(\phi_\theta^T(\hat{x})\mathbf{u}_w)N(\mathbf{w} | \mathbf{u}_w, \Sigma_w)d\mathbf{w} \approx \sigma(\kappa(\sigma_x^2)\mathbf{u}_w^T \phi_\theta(\hat{x})),$$

where $\kappa(\sigma_x^2) = (1 + \frac{\pi}{8}\phi_\theta^T(\hat{x})\Sigma_w\phi_\theta(\hat{x}))^{-1/2}$ is the variance of \hat{x} with the covariance of sample posterior distributions Σ_w .

To arrive at a binary classification, we choose $\mathbf{u}_w^T \phi_\theta(\hat{x}) = 0$ as the decision boundary, where we have the probability $p(\hat{y} = 1 | \hat{x}, \mathbf{S}) = 0.5$. Thus, computing the sign of $\mathbf{u}_w^T \phi_\theta(\hat{x})$ will meet the case of 0-1 classification. Moreover, the likelihood of prediction provides the confidence of the prediction, which is more important in unbalanced classification tasks.

D. Implementation

We detail the implementation of PFCVM_{LP} step by step and provide pseudo-code in Algorithm 1.

Algorithm 1 PFCVM_{LP} algorithm

```

1: Input: Training data set:  $\mathbf{S}$ ;
   initial values:  $INITVALUES$ ;
   threshold:  $THRESHOLD$ ;
   maximum number of iterations:  $maxIts$ .
2: Output: Weights of model:  $WEIGHT$ ;
   hyperparameters:  $HYPERPARAMETER$ .
3: Initialization:  $[\mathbf{w}, \theta, \alpha, \beta] = INITVALUES$ 
4: Index = generateIndex( $\alpha, \beta$ )
5: while  $i < maxIts$  do
6:    $\Phi = updateBasisFunction(\mathbf{x}, \theta, \mathbf{Index})$ 
7:    $[\mathbf{w}, \theta] = updatePosterior(\Phi, \mathbf{w}, \theta, \alpha, \beta, \mathbf{Y})$ 
8:    $[\alpha, \beta] = maximumMarginal(\Phi, \mathbf{w}, \theta, \alpha, \beta, \mathbf{Y})$ 
9:   if  $\alpha_i$  or  $\beta_k > THRESHOLD.maximum$  then
10:    delete the  $i$ th sample or the  $k$ th feature
11:   end if
12:   Index = updateIndex( $\alpha, \beta$ )
13:    $marginal = calculateMarginal(\Phi, \mathbf{w}, \theta, \alpha, \beta, \mathbf{Y})$ 
14:   if  $\Delta marginal < THRESHOLD.minimal$  then
15:     break
16:   end if
17:    $WEIGHT = [\mathbf{w}, \theta, \mathbf{Index}]$ 
18:    $HYPERPARAMETER = [\alpha, \beta]$ 
19: end while

```

Algorithm 1 consists of the following main steps.

- 1) First, the values of $\mathbf{w}, \theta, \alpha, \beta$ are initialized by $INITVALUES$ and a parameter **Index** generated to indicate the useful samples and features (line 3 and 4).
- 2) At the beginning of each iteration, compute the matrix Φ according to (3) (line 6).

- 3) Based on (9), use the new hyperparameters to re-estimate the posterior (line 7).
- 4) Use the re-estimated parameters to maximize the logarithm of marginal likelihood according to (14) (line 8).
- 5) Prune irrelevant samples and useless features if the corresponding hyperparameters are larger than a specified threshold (lines 9, 10, 11).
- 6) Update the **Index** vector (line 12).
- 7) Calculate the logarithm of the marginal likelihood (line 13).
- 8) Convergence detection, if the change of marginal likelihood is relatively small, halt the iteration (line 14, 15, 16).
- 9) Generate the output values. The vector **WEIGHT** consists of sample and feature weights and the vector **Index** indicates the relevant samples and features (line 17, 18).

We have now presented all details of PFCVM_{LP} , including derivations of equations and pseudo-code. Next, we evaluate the performance of PFCVM_{LP} by comparisons with other state-of-the-art algorithms on a synthetic dataset, low-dimensional benchmark datasets, and high-dimensional gene expression datasets.

IV. EXPERIMENTAL RESULTS

In a series of experiments we assess the performance of PFCVM_{LP} . The first experiment aims to robustness the ability of PFCVM_{LP} against noisy features. Second, a set of experiments are carried out on 11 benchmark datasets to assess the performance of classification. Then, experiments are designed on gene expression datasets, which contain lots of irrelevant features. Finally, the computational and space complexity of PFCVM_{LP} is analyzed.

A. Synthetic Dataset: Robustness Against Noise

The first experiment is aimed at verifying the robustness of PFCVM_{LP} against increasing numbers of irrelevant features. We choose RVM,⁷ SVM (LIBSVM⁸), WSVM⁹ and JCFO [12] as control groups. We randomly draw 200 training samples and 500 testing samples for each class to synthesize this dataset. Using the procedure defined in [12], each sample is drawn from one of two different Gaussian distributions, which share the identical unit standard deviations but opposite means. The means of the distributions in the r -dimensional space are given by:

$$\mu_1 = -\mu_2 = [1/\sqrt{2}, 1/\sqrt{2}, \underbrace{0, 0, \dots, 0}_{r-2}],$$

where r stands for the dimension of synthetic data. So, for each dataset, there are $r - 2$ irrelevant features.

Regardless of irrelevant features, the optimal classifier should be a line across the original point with the slope of -1 with the Bayesian error rate $\Phi(-1) \approx 0.1597$. To make the results more convincing, we randomly sample data samples to




Fig. 2. Performance of RVM, SVM, WSVM, JCFO and PFCVM_{LP} with increasing numbers of irrelevant features. The X-axis is the number of features and the Y-axis is the accuracy. For each dataset, there are only 2 useful features.

generate 20 distinct training and testing sets, then we run each algorithm, on each data partition. We choose a linear kernel as the basis function. The parameters of the rival methods are chosen by cross-validation. The results, shown in Fig. 2, are the average of 20 results. Clearly, PFCVM_{LP} and JCFO outperform the other methods when $r > 20$.

These series of experiments demonstrate that in the presence of irrelevant features, PFCVM_{LP} achieves comparable performance to JCFO, and performs better than other rival algorithms. PFCVM_{LP} , JCFO and WSVM are classification and feature selection co-training algorithms; when $r \leq 20$, they are largely immune to irrelevant features, however, as we continue to add irrelevant features, WSVM is not able to maintain this immunity while PFCVM_{LP} and JCFO are only slightly effected. Only PFCVM_{LP} and JCFO are able to select relevant features since they can estimate the contributions of input features by using feature parameters. The rapid degradation of the performance of RVM and SVM demonstrates the necessity of feature selection.

B. Benchmark Datasets: Classification Performance

The benchmark datasets consist of 1 synthetic dataset, banana, and 10 real-world datasets from DELVE¹⁰ and UCI [21]. Information on these datasets is reported in Table I. These datasets have relatively low-dimensional features and we could hardly assess which of them are relevant features and which of them are not. So we design this series of experiments to evaluate the classification performance and probabilistic predictions of PFCVM_{LP} especially when the feature selection is sometimes not crucial.

Some of datasets have been organized by Rätsch et al. [22]. For the others, following the implementation reported by Chen et al. [8], we convert them into binary classification cases and randomly split each dataset into 100 partitions. Then we

⁷<http://www.miketipping.com/>

⁸<https://www.csie.ntu.edu.tw/~cjlin/libsvm/>

⁹http://www.humansensing.cs.cmu.edu/down_svmfeat_main.php

¹⁰<http://www.cs.toronto.edu/~delve/data/datasets.html>

TABLE I
BASIC INFORMATION ABOUT OUR BENCHMARK DATASETS. “POSITIVE”
INDICATES THE RELATIVE SIZE OF THE POSITIVE CLASS, I.E., $y = +1$.

| Dataset | # Train | # Test | Positive | # Features |
|------------|---------|--------|----------|------------|
| Australian | 414 | 276 | 44.49% | 14 |
| Banana | 400 | 4900 | 44.83% | 2 |
| Breast | 200 | 77 | 29.28% | 9 |
| Diabetics | 468 | 300 | 34.90% | 8 |
| Image | 1300 | 1010 | 56.95% | 18 |
| Ionosphere | 210 | 141 | 64.10% | 34 |
| Solar | 666 | 400 | 65.28% | 9 |
| Splice | 1000 | 2175 | 44.93% | 60 |
| Thyroid | 140 | 75 | 30.23% | 5 |
| Titanic | 150 | 2051 | 58.33% | 3 |
| Waveform | 400 | 4600 | 32.94% | 21 |

preprocess these datasets before performing experiments. In our preprocessing procedure, we have two normalization steps: sample-wise we normalize each dataset to have mean 0 and variance 1; dimension-wise we then normalize them to have mean 0 and variance 1.

In order to assess the classification performance, we compare our algorithms with RVM, PCVM, soft-margin SVM (SVM_{soft}) [23], $PCVM_{mRMR}$, $PFCVM_{EM}$, WSVM and sparse multinomial logistic regression via Bayesian L1 regularization (SMLRB) [24]. The features for $PCVM_{mRMR}$ are selected using the minimal redundancy maximal relevance criterion (mRMR) [14] in advance. Among the algorithms considered, mRMR, SMLRB, WSVM and $PFCVM_{EM}$ are feature selection algorithms. We choose k nearest neighbors (k -NN), where k is determined by cross-validation from $\{1, 2, \dots, 20\}$ and linear discriminant analysis (LDA) as baseline learning algorithms. In these experiments, the Gaussian RBF is employed as the basis function. Two popular evaluation criteria, i.e., the area under the curve of receiver operating characteristic (AUC) and error rate (ERR)¹¹ are used as for evaluation; they represent a probability criterion and a threshold criterion, respectively [25].

We follow the procedure in [8] to choose optimal initial model parameters, in which the first five partitions of each dataset are chosen for cross-validation: we train each algorithm with all parameter candidates and then choose the parameters with the lowest median error rate on these five partitions.

We use this procedure to choose the optimal k for k -NN and the kernel parameter ϑ for RVM, PCVM and WSVM. For SVM_{soft} , the regularization C and the kernel parameter ϑ are tuned by grid search, in which we train SVM_{soft} with all combinations of each candidate C and ϑ , then choose the combination with the lowest median error rate. For $PCVM_{mRMR}$, the proper sizes of feature subsets and the kernel parameter ϑ for PCVM are chosen by the similar grid search. We also choose the proper starting points for $PFCVM_{EM}$ and the initial hyperparameters for $PFCVM_{LP}$.

Results of the experiments are reported in Table II. We also report results of paired t-tests of $PFCVM_{LP}$ versus other rival algorithms, which assess whether the distributions of error

rate or AUC of two classifiers are statistically different (with p -value < 0.05). In terms of ERR, $PFCVM_{LP}$ outperforms other methods in 5 out of 11 cases. In terms of AUC, $PFCVM_{LP}$ also beats other methods in 5 cases. In terms of either ERR or AUC, SVM_{soft} only wins on 2 datasets. RVM, PCVM and $PFCVM_{EM}$ perform similarly. Overall, $PFCVM_{LP}$ is the most frequent winner.

Compared to PCVM, in terms of accuracy $PFCVM_{LP}$ beats PCVM on 9 datasets and in terms of AUC $PFCVM_{LP}$ wins on 7 datasets. These results demonstrate that simultaneously learning feature weights increases the performance of classification. Comparing with $PFCVM_{EM}$, we note that $PFCVM_{LP}$ wins on 10 datasets (accuracy) and 9 datasets (AUC). Moreover, averaged across all benchmark datasets $PFCVM_{LP}$ selects 7.96 features, while $PFCVM_{EM}$ selects 10.25. With these results we conclude that $PFCVM_{LP}$ yields more accurate and sparser classifiers than $PFCVM_{EM}$. Overall, complete Bayesian estimation outperforms the EM algorithm based MAP point estimation, which can only guarantee a locally optimal solution and is sensitive to initial values.

C. High-Dimension Gene Expression Data: Performance in the Presence of Many Irrelevant Features

Gene expression datasets contain lots of irrelevant features, which degrades the performance of classifiers [26–28]. In our third set of experiments, we examine whether $PFCVM_{LP}$ is able to eliminate the irrelevant features and make informative predictions.

We choose three popular gene expression datasets:¹² colon cancer [26], Duke cancer [29] and leukemia [27]. The colon cancer dataset includes 2000 features from 22 normal and 40 colon tissues. The Duke cancer dataset is from Duke Breast Cancer SPORE frozen tissue bank, and consists of 42 instances with 7129 features. Each instance can be separated by the states of the estrogen and progesterone receptors. The leukemia dataset has 7129 features from 25 acute lymphoblastic leukemia (ALL) and 47 acute myeloid leukemia (AML) patients.

Considering the relatively small numbers of samples versus the large numbers of features, in this set of experiment we use the leave one out cross validation (LOOCV) method to conduct the experiments: each time a sample is left out to be diagnosed, and the classification model is trained to fit the remaining data. So, we generate 62, 42 and 72 runs for the colon cancer dataset, the Duke cancer dataset and the leukemia dataset, respectively. Following [30], in preprocessing each dataset is normalized in two ways: sample-wise to follow a standard normal distribution and then dimension-wise to follow a standard normal distribution.

We compare $PFCVM_{LP}$ with typical feature selection algorithms, including the minimal redundancy maximal relevance criterion (mRMR) [14], the Laplacian score (LS) [16] and Student’s t-test [15], with PCVM as the classifier. As before, we also compare $PFCVM_{LP}$ with other feature and classifier co-learning algorithms, i.e., $PFCVM_{EM}$ and JCFO. Moreover,

¹¹ERR = $\frac{1}{N} \sum_{i=1}^N 1(y_i \neq f(\mathbf{x}_i; \mathbf{w}, \theta))$

¹²These datasets are available from <http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>

TABLE II

RESULTS OF EXPERIMENTS ON BENCHMARK DATASETS. RESULTS MARKED WITH A SUPERScript * INDICATE THAT THE DISTRIBUTION OF THE RESULT IS STATISTICALLY DIFFERENT FROM $PFCVM_{LP}$. FOR LDA WE FAIL TO OBTAIN RESULTS ON THE IONOSPHERE DATASET, BECAUSE THE POOLED COVARIANCE MATRIX IS NOT POSITIVE DEFINITE.

| Error | Australian | Banana | Breast | Diabetics | Image | Ionosphere | Solar | Splice | Thyroid | Titanic | Waveform |
|--------------|----------------|----------------|----------------|--------------|-------------|-------------|--------------|-------------|-------------|--------------|----------------|
| k -NN | 14.14* | 11.26* | 29.98* | 25.83* | 2.82* | 14.51* | 35.83* | 22.02* | 4.36* | 30.08* | 11.00 |
| LDA | 13.67* | 47.00* | 27.17 | 23.41 | 16.89* | — | 34.41 | 16.15* | 14.44* | 22.67 | 14.94* |
| SVM | 15.62* | 10.42 * | 28.94* | 23.43* | 3.27* | 8.45* | 35.48* | 10.46* | 4.84* | 22.45 | 10.04 * |
| RVM | 13.66* | 10.70* | 27.55* | 23.73 | 4.08* | 8.10 | 33.66 | 12.77* | 4.21* | 22.89 | 10.86* |
| PCVM | 13.68* | 10.67* | 27.09 | 23.74 | 4.91* | 11.84* | 34.75* | 13.57* | 4.64* | 22.58* | 10.65 |
| mRMR | 13.37 * | 10.57 | 26.45 | 23.73 | 3.29* | 9.78* | 34.09* | 9.83* | 5.35* | 39.60* | 10.65* |
| SMLRB | 13.71* | 55.11* | 32.27* | 26.46* | 17.76* | 19.77* | 33.04 | 16.26* | 10.12* | 23.72* | 22.57* |
| WSVM | 14.13* | 33.54* | 29.66* | 23.51 | 5.55* | 8.35 | 34.75* | 16.03* | 6.28* | 22.65 | 13.06* |
| $PFCVM_{EM}$ | 15.03* | 10.80* | 25.48 * | 25.36* | 3.29 | 8.77* | 34.97* | 7.88 | 3.87 | 22.95 | 11.15* |
| $PFCVM_{LP}$ | 14.95 | 10.64 | 26.68 | 23.67 | 1.97 | 7.82 | 33.65 | 7.12 | 3.59 | 22.00 | 11.10 |

| AUC | Australian | Banana | Breast | Diabetics | Image | Ionosphere | Solar | Splice | Thyroid | Titanic | Waveform |
|--------------|----------------|--------------|----------------|----------------|--------------|----------------|--------------|--------------|--------------|--------------|--------------|
| k -NN | 85.85* | 88.40* | 60.12* | 68.00* | 96.97* | 81.29* | 64.30* | 77.31* | 94.14* | 63.36* | 88.63* |
| LDA | 92.88 * | 53.31 | 71.19* | 82.71* | 81.80* | — | 66.53* | 91.88* | 86.67* | 71.09* | 86.90* |
| SVM | 89.95* | 96.15 | 64.87* | 82.94 * | 99.37* | 97.57 * | 68.69* | 96.13 | 98.99* | 70.70* | 96.46 |
| RVM | 92.44* | 96.20 | 71.68* | 82.76* | 96.44* | 95.35 | 74.06* | 94.09* | 99.01 | 73.03* | 96.09* |
| PCVM | 92.37* | 96.10 | 71.84 * | 82.76* | 98.98* | 93.72* | 72.05* | 93.66* | 98.32* | 75.42* | 96.11* |
| mRMR | 92.32* | 96.19 | 69.62* | 82.06* | 99.27* | 93.89* | 72.90 | 95.88* | 97.99* | 68.20* | 96.07* |
| SMLRB | 86.61* | 51.53* | 66.35 | 74.76* | 81.09* | 73.13* | 66.98* | 83.70* | 84.70* | 69.95* | 83.10* |
| WSVM | 86.54* | 67.99* | 62.69* | 82.59* | 98.57* | 96.23* | 67.43* | 91.73* | 97.42* | 69.71* | 93.96* |
| $PFCVM_{EM}$ | 90.80 | 96.07 | 68.66* | 78.98* | 99.45* | 95.43 | 72.08* | 95.67* | 98.79* | 73.43* | 94.88* |
| $PFCVM_{LP}$ | 90.78 | 96.12 | 68.18 | 79.98 | 99.72 | 95.35 | 74.08 | 96.94 | 99.07 | 75.50 | 95.44 |




Fig. 3. Accuracy curves of different algorithms with different scales of selected features. LS, T-test and mRMR fail to select informative features on the leukemia dataset.

in order to validate whether the selected subsets of features by $PFCVM_{LP}$ also work for other classifiers, we also conduct experiments with RVM, SVM (LIBSVM), RVM_{PFCVM} and SVM_{PFCVM} , whose features are selected by $PFCVM_{LP}$, in advance. As the dimensions of gene data are relatively high, we choose the inner product, i.e., the linear kernel, as the metric. Finally, we report the averages of all runs on each dataset as the results, i.e., the averages of 62, 42 and 72 runs on the colon cancer, Duke cancer and leukemia datasets, respectively.

Fig. 3 shows the accuracy curves of feature selection algorithms. For mRMR, LS and T-test, the span of the selected features is $[10, 20, \dots, 80]$. We train a SVM with each number of selected features, and then plot the accuracy. Comparing our algorithms with others, from Fig. 3 we see that $PFCVM_{LP}$ achieves the highest accuracy scores on 2 out of 3 datasets. Moreover, on average $PFCVM_{LP}$ selects the smallest subsets

of features on every dataset. However, the performance of $PFCVM_{EM}$ is also competitive. In Fig. 4, we illustrate the cumulative number¹³ of occurrences for each feature on each dataset with $PFCVM_{LP}$ and $PFCVM_{EM}$. The number of features selected by the two algorithms are similar on average, but from Fig. 4, we see that $PFCVM_{LP}$ concentrates on smaller sets of relevant features. This result shows that a complete Bayesian solution is more stable than a solution based on the EM algorithm.

On the colon cancer dataset, $PFCVM_{LP}$ selects 4.94 features, on average. Among all 2000 genes, 5 of them are particularly important (occurring in more than half of the tests). The biological explanations of these 5 genes are reported in Table III and 2 of them (No. 377 and No. 1772) are the same

¹³We use the phrase “cumulative number” to refer to the number of times a feature is selected in the experiments, e.g., if a feature is never selected in the experiments, its “cumulative number” is 0.

genes selected by Li et al. [31]. On the Duke cancer dataset, $PFCVM_{LP}$ on average selects 2.14 features and 2 of them are selected in almost every run. On the leukemia dataset, $PFCVM_{LP}$ select 2.94 features on average, and 5 of the 7 most occurred genes are among the 50 genes most correlated with the diagnosis [27]. The biological significance of these genes are reported in Table IV.

TABLE III
BIOLOGICAL SIGNIFICANCE OF THE MOST FREQUENTLY OCCURRING GENES IN THE COLON CANCER DATASET. # DENOTES THE NUMBER OF OCCURRENCES FOR A FEATURE IN ALL RUNS.

| Data ID | # | GenBank ID | Description [26] |
|---------|----|------------|------------------------------------|
| 1772 | 61 | 0H8393 | Collagen alpha 2(XI) chain |
| 1668 | 60 | M82919 | mRNA for GABAA receptor |
| 1210 | 58 | R55310 | Mitochondrial processing peptidase |
| 377 | 51 | R39681 | Eukaryotic initiation factor |
| 1679 | 37 | X53586 | mRNA for integrin alpha 6 |

TABLE IV
BIOLOGICAL SIGNIFICANCE OF THE MOST FREQUENTLY OCCURRING GENES IN THE LEUKEMIA DATASET. # DENOTES THE NUMBER OF OCCURRENCES FOR A FEATURE IN ALL RUNS. THE SUPERScript * DENOTES THAT THESE GENES ARE AMONG THE TOP 50 MOST IMPORTANT GENES FOR DIAGNOSING AML/ALL [27].

| Data ID | # | GenBank ID | Relation | Description [27] |
|---------|----|------------|----------|------------------------|
| 4847* | 70 | X95735 | AML | Zyxin |
| 4951 | 68 | Y07604 | N/A | Nucleoside diphosphate |
| 6169 | 11 | M13690 | AML | Hereditary angioedema |
| 3847* | 6 | U82759 | AML | HoxA9 mRNA |
| 2354* | 6 | M92287 | AML | CCND3, Cyclin D3 |
| 4973* | 5 | Y08612 | ALL | Protein RABAPTIN-5 |
| 1834* | 5 | M23197 | AML | CD33 antigen |

Table V reports the results of RVM, SVM, RVM_{PFCVM} , SVM_{PFCVM} and $PFCVM_{LP}$. According to this table, with the selected subsets of features, RVM_{PFCVM} and SVM_{PFCVM} outperform the original methods RVM and SVM, respectively. This comparison shows that the selected subsets of features also work well for other algorithms, like RVM and SVM. Furthermore, SVM_{PFCVM} even beats $PFCVM_{LP}$ and outputs the best prediction on the Duke cancer dataset. This result is reasonable, because in most cases, as shown in Fig. 4, $PFCVM_{LP}$ only selects two features to make decisions and these two features are enough for SVM to make accurate decisions, shown in Fig. 5.

TABLE V
ACCURACY OF DIAGNOSES ON GENE EXPRESSION DATASETS.

| Accuracy (%) | Colon Cancer | Duke Cancer | Leukemia |
|---------------|--------------|--------------|--------------|
| RVM | 85.48 | 80.95 | 93.06 |
| SVM | 83.87 | 85.71 | 87.50 |
| RVM_{PFCVM} | 87.10 | 92.86 | 95.83 |
| SVM_{PFCVM} | 85.48 | 97.62 | 95.83 |
| $PFCVM_{LP}$ | 96.77 | 95.24 | 98.61 |

These experiments demonstrate that $PFCVM_{LP}$ cannot only automatically determine a proper subset of informative features but also make reliable predictions. The selected features also work well for other classifiers, like RVM and SVM.




Fig. 5. Duke cancer data in the selected feature subspace. In this subspace, the two classes could be separated by a line, e.g., the dash is one of the proper decision boundary between two classes.

D. Complexity Analysis

While computing the posterior covariance Σ_θ in (11), we have to derive the negative inverse of the Hessian matrix. This derivation does not guarantee a numerically accurate result, because of the ill-condition of this Hessian matrix. Practically, we abandon the term \mathbf{E} in (11), so that the calculation becomes:

$$\Sigma_\theta = (\mathbf{D}^T \mathbf{C} \mathbf{D} + \mathbf{B} + \mathbf{O}_\theta)^{-1}. \quad (16)$$

In (16), \mathbf{B} and \mathbf{O}_θ are positive definite diagonal matrices and $\mathbf{D}^T \mathbf{C} \mathbf{D}$ has a quadratic form. Theoretically, the Hessian is a positive definite matrix. Nevertheless, because of machine precision, ill-condition may still occur occasionally, especially when β_k is very large.

In the case of large β_k , especially when $\beta_k \rightarrow \infty$, the corresponding feature weight θ_k is restricted to a small neighborhood around 0. So, during the iteration, we delete this feature weight θ_k and this hyperparameter β_k from our model. Initially, all the features are contained in the model. The main computational cost is the Cholesky decomposition in computing covariances of posteriors, Σ_θ and Σ_w , which is $O(N^3 + M^3)$. Thus the computational complexity of $PFCVM_{LP}$ is the same as for RSFM [18].

As for the storage requirements for $PFCVM_{LP}$, the basis function matrix Φ_θ needs $O(N^2)$ for storage, and in the initial training stage the covariance matrices Σ_θ and Σ_w require $O(M^2)$ and $O(N^2)$ storage, respectively. Therefore, the overall space complexity of $PFCVM_{LP}$ is $O(N^2 + M^2)$, which is slightly better than RSFM with a $O(N^2 M + M^2)$ space complexity. As iterations proceed, N and M are rapidly decreasing, resulting in $O(\bar{N}^3 + \bar{M}^3)$ computational complexity and $O(\bar{N}^2 + \bar{M}^2)$ space complexity, where $\bar{N} \ll N$ and $\bar{M} \ll M$. In our experiments, N and M rapidly decrease to relatively small numbers in the first few iterations and the training speed quickly accelerates.

As illustrated by Fig. 6, during the first 40 iterations the size of features and samples decreases from 60 to 17 and from 1000 to 127, respectively, and the CPU time for each iteration step is decreased to 2.7% of the first iteration.




Fig. 4. Illustrations of selected features on gene expression datasets. The horizontal axis shows the index of features and the vertical axis shows the cumulative number of occurrences for the corresponding feature. The dashed line in each figure indicates the maximum cumulative number. The figures at the top show the features selected by PFCVM_{LP} ; those at the bottom show the features selected by PFCVM_{EM} .




Fig. 6. Illustration of the rapid decrease in the size of features, samples and CPU time. In these experiments, we choose the splice dataset as an example.

V. GENERALIZATION AND SPARSITY

Both the benchmark and gene expression experiments indicate that the proposed classifier and feature selection co-learning algorithm is capable of generating a sparse solution. In this section, we first analyse the KL-divergence between the prior and posterior. Following this, we investigate the entropic constraint Rademacher complexity [32] and derive a generalization bound for PFCVM_{LP} . By tightening the bound, we theoretically demonstrate the significance of sparsity assumption and introduce a method to choose the initial values for PFCVM_{LP} .

A. KL-Divergence Between Prior and Posterior

In Bayesian learning, we use KL-Divergence to measure the information gain from prior to posterior. As discussed in Section III-A, the approximated posterior over feature parameters, denoted as $\tilde{q}(\theta) = \mathcal{N}(\theta | \mathbf{u}_\theta, \Sigma_\theta)$, is a multivariate Gaussian distribution. However, as the feature prior is the left-truncated Gaussian prior, the true posterior over feature parameters

should be restricted to the positive quadrant. In order to achieve this we first compute the probability mass of the posterior in this half, $Z_0 = \int_0^\infty \tilde{q}(\theta) d\theta$; after that we obtain a re-normalized version of the posterior: $q(\theta) = \tilde{q}(\theta)/Z_0$, where $\theta_k \geq 0$.

We denote $\beta_0 = (\beta_{0,1}, \beta_{0,2}, \dots, \beta_{0,M})$ as the initial prior and $\beta = (\beta_1, \beta_2, \dots, \beta_M)$ as the optimized prior. Following [3], we adopt the independent posterior assumption. We compute $KL_\theta(q||p)$ ¹⁴ using the following formula (the details

¹⁴ $KL_\theta(q||p)$ denotes the KL-divergence between the posterior and prior in feature weights; p and q are short for $p(\theta | \beta)$ and $q(\theta)$, respectively.

are specified by [33]):

$$\begin{aligned} KL_{\theta}(q||p) &= \int_0^{\infty} q(\theta) \frac{\ln q(\theta)}{\ln p(\theta | \beta_0)} d\theta \\ &= \sum_{k, \theta_k \neq 0} \left\{ \begin{aligned} &\frac{1}{2} \left[\frac{\beta_{0,k}}{\beta_k} - 1 + \ln \left(\frac{\beta_k}{\beta_{0,k}} \right) + \beta_{0,k} \theta_k^2 \right] \\ &+ \frac{(2\pi\beta_k)^{-1/2} (\beta_{0,k} + \beta_k) \theta_k}{\operatorname{erfcx}(-\theta_k \sqrt{\beta_k/2})} Z_{0,k}^{-1} \\ &- \ln \left(\operatorname{erfc} \left(-\frac{\theta_k \beta_k}{2} \right) \right) \end{aligned} \right\}, \end{aligned}$$

where $\operatorname{erfc}(z) = \frac{2}{\sqrt{\pi}} \int_z^{\infty} e^{-t^2} dt$ and $\operatorname{erfcx}(z) = e^{z^2} \operatorname{erfc}(z)$.

Note that $Z_{0,k} = \int_0^{\infty} \tilde{q}(\theta_k) d\theta_k = 1/2 \operatorname{erfc}(-\theta_k \sqrt{\beta_k/2})$, so we can calculate $KL_{\theta}(q||p)$ as:

$$\begin{aligned} KL_{\theta}(q||p) &= \sum_{k, \theta_k \neq 0} \left\{ \begin{aligned} &\frac{1}{2} \left[\frac{\beta_{0,k}}{\beta_k} - 1 + \ln \left(\frac{\beta_k}{\beta_{0,k}} \right) + \beta_{0,k} \theta_k^2 \right] \\ &+ \frac{(2\pi\beta_k)^{-1/2} (\beta_{0,k} + \beta_k) \theta_k}{2 \exp(\beta_k \theta_k^2/2)} Z_{0,k}^{-1} \\ &- \ln(Z_{0,k}) + \ln \left(\frac{\operatorname{erfc}(-\theta_k \sqrt{\beta_k/2})}{2 \operatorname{erfc}(-\theta_k \beta_k/2)} \right) \end{aligned} \right\}. \end{aligned}$$

The KL-divergence is dominated by two parameters: θ and β . However, the sensitivity of $KL_{\theta}(q||p)$ to these two parameters is different. As shown in Fig. 7, setting the initial hyperparameter $\beta_{0,k} = 0.5$, we see that when changing the value of θ , the curve of $KL_{\theta}(q||p)$ shows significant changes, while this curve changes little when changing the optimized β_k . Also, the minimum of $KL_{\theta}(q||p)$ is near the $\theta_k = 0$ axis, where the corresponding feature is pruned.




Fig. 7. Illustration of the numerical contribution of θ and β to $KL_{\theta}(q||p)$. To evaluate this contribution alone, we assume each sample weight $w_i = 1$ and each sample hyperparameter $\alpha_i = 1$. (Best viewed in color.)

B. Rademacher Complexity Bound

For a binary classification learning problem the goal is to learn a function $f : \mathbb{R}^M \rightarrow \{-1, +1\}$ from a hypothesis class F , with the given dataset $\mathbf{S} = \{x_i, y_i\}_{i=1}^N$ drawn i.i.d. from distribution D . We attempt to assess f by the expectation loss: $L(f) = E_{(x,y) \sim D} l(y, f(x))$, where $l(y, f(x))$ is a loss function. Practically, D is inaccessible and we can only assess the empirical loss for the given dataset \mathbf{S} : $\Lambda(f, \mathbf{S}) = \frac{1}{N} \sum_{i=1}^N l(y_i, f(x_i))$. We adopt a 0-1

loss function: $l_{0-1}(y, f(x)) = I(yf(x) \geq 0)$, where $I(\cdot)$ is the indicator function. The loss function is dominated by the $1/c$ -Lipschitz function: $l_c(a) = \min(1, \max(0, 1 - a/c))$, namely $l_{0-1}(y, f(x)) \leq l_c(yf(x))$. Then, we conclude the entropic constraint Rademacher complexity bound in the following theorem:

Theorem 1: [3, 32] Based on the posterior $q(\mathbf{w}, \theta)$ given in Section III-A, we have the Bayes voting classifier:

$$\hat{y} = f(x, q) = E_{q(\mathbf{w}, \theta)} [\operatorname{sign}(\Phi_{\theta}(\mathbf{x})\mathbf{w})]. \quad (17)$$

Define $r > 0$ and $g > 0$ as arbitrary parameters. For all $f \in F$, defined in the start of Section V-B, with probability at least $1 - \sigma$, the following bound for the generalization error on a given dataset \mathbf{S} holds:

$$\begin{aligned} P(yf(x) < 0) &\leq \Lambda(f, \mathbf{S}) + \frac{2}{c} \sqrt{\frac{2\tilde{g}(q(\mathbf{w}, \theta))}{N}} \\ &\quad + \sqrt{\frac{\ln \log_r \frac{r\tilde{g}(q(\mathbf{w}, \theta))}{g} + \frac{1}{2} \ln \frac{1}{\delta}}{N}}, \end{aligned}$$

where c is the $1/c$ -Lipschitz parameter, the empirical loss $\Lambda(f, \mathbf{S}) = \frac{1}{N} \sum_{i=1}^N l_c(y_i f(x_i))$ and the Rademacher entropic constraints $\tilde{g}(q(\mathbf{w}, \theta)) = r \cdot \max\{KL(Q||P), g\}$. $KL(Q||P)$ is the KL-divergence between the posteriors and priors in sample and feature weights.

With a constant training set, the expectation loss is mainly bounded by the empirical loss and $KL(Q||P)$. When the empirical loss is acceptable, a smaller $KL(Q||P)$ could decrease the complexity term and leads to a tighter bound. The contribution of \mathbf{w} (the sample weight) has been analysed in [3]. In order to analyse the effect of θ (the feature weight) alone, we can assume the \mathbf{w} is a given constant value $\hat{\mathbf{w}}$. As a result, we have $q(\hat{\mathbf{w}}, \theta) = q(\theta | \hat{\mathbf{w}})$ and thus $KL(Q||P) = KL_{\theta|\hat{\mathbf{w}}}(q||p)$.

As shown in Fig. 7, the minimal $KL_{\theta|\hat{\mathbf{w}}}(q||p)$ is near $\theta_k = 0$. This is consistent with our prior assumption and demonstrates that a truncated Gaussian prior over features can benefit the generalization performance by running as a regularization term and simultaneously encourage sparsity in feature space. Furthermore, in our model, the posterior and marginal likelihood are maximized iteratively in the training step. To accelerate the speed of convergence, we may choose proper starting points by minimizing $KL_{\theta|\hat{\mathbf{w}}}(q||p)$, i.e., as indicated in the end of Section V-A, we can use an optimal β instead of β_0 as initial hyperparameter.

VI. CONCLUSION

We have proposed a joint classification and feature learning algorithm PFCVM_{LP}. The proposed algorithm adopts sparseness-promoting priors for both sample and feature weights to jointly learn to select relevant samples and features. By using the Laplace approximation, we compute a complete Bayesian estimation of PFCVM_{LP}, which is more stable than previously considered EM-based solutions. The performance of PFCVM_{LP} has been examined according to two criteria: the accuracy of its classification results and its ability to select informative features. Our experiments demonstrate that the accuracy of PFCVM_{LP} on benchmark data sets is either the

best or close to the best. On high-dimensional gene expression datasets, PFCVM_{LP} performs highly accurately when compared to other approaches. Finally, a Rademacher complexity bound is derived for the proposed method. By tightening this bound, we demonstrate the significance of feature selection and introduce a way of finding proper initial values.

PFCVM_{LP} jointly encourages sparsity to features and samples. However, in order to select features for non-linear basis functions, we have to differentiate, which leads to high computational costs. As future work, we plan to use incremental learning [3, 34] to reduce the computational costs. We also plan to design an online strategy [35] for joint feature and classifier learning. Also, PFCVM_{LP} focuses on the binary classification problem. It would be interesting to extend PFCVM_{LP} to multiple classes [36, 37]. Finally, we aim to use PFCVM_{LP} in other areas of research, such as in bioinformatics problems and clinical diagnoses [38].

REFERENCES

- [1] C. Cortes and V. Vapnik, "Support-vector networks," *Machine learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [2] C. Li and H. Chen, "Sparse Bayesian approach for feature selection," in *Proceedings of IEEE Symposium on Computational Intelligence in Big Data (CIBD), Orlando, FL, USA, December 9-12, 2014*, pp. 7–13. [Online]. Available: <http://dx.doi.org/10.1109/CIBD.2014.7011521>
- [3] H. Chen, P. Tino, and X. Yao, "Efficient probabilistic classification vector machine with incremental basis function selection," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 25, no. 2, pp. 356 – 369, 2014.
- [4] S. He, H. Chen, Z. Zhu, D. G. Ward, H. J. Cooper, M. R. Viant, J. K. Heath, and X. Yao, "Robust twin boosting for feature selection from high-dimensional omics data with label noise," *Information Sciences*, vol. 291, pp. 1–18, 2015.
- [5] Y. Mohsenzadeh, H. Sheikhzadeh, and S. Nazari, "Incremental relevance sample-feature machine: A fast marginal likelihood maximization approach for joint feature selection and classification," *Pattern Recognition*, vol. 60, pp. 835–848, 2016.
- [6] V. N. Vapnik, *Statistical Learning Theory*. Wiley New York, 1998, vol. 2.
- [7] M. E. Tipping, "Sparse Bayesian learning and the relevance vector machine," *The Journal of Machine Learning Research*, vol. 1, pp. 211–244, 2001.
- [8] H. Chen, P. Tino, and X. Yao, "Probabilistic classification vector machines," *IEEE Transactions on Neural Networks*, vol. 20, no. 6, pp. 901–914, 2009.
- [9] J. O. Berger, *Statistical Decision Theory and Bayesian Analysis*, 2nd ed. Springer, 1985.
- [10] C.-C. Chang and C.-J. Lin, "LIBSVM: a library for support vector machines," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 2, no. 3, p. 27, 2011.
- [11] R. E. Bellman, *Adaptive Control Processes: A Guided Tour*. Princeton University Press, Princeton, 1961.
- [12] B. Krishnapuram, A. Hartemink, L. Carin, and M. A. Figueiredo, "A Bayesian approach to joint feature selection and classifier design," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 9, pp. 1105–1111, 2004.
- [13] M. H. Nguyen and F. De la Torre, "Optimal feature selection for support vector machines," *Pattern Recognition*, vol. 43, no. 3, pp. 584–591, 2010.
- [14] H. Peng, F. Long, and C. Ding, "Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 8, pp. 1226–1238, 2005.
- [15] D. C. Montgomery, G. C. Runger, and N. F. Hubele, *Engineering statistics*. John Wiley & Sons, 2009.
- [16] X. He, D. Cai, and P. Niyogi, "Laplacian score for feature selection," in *Advances in Neural Information Processing Systems 18, Papers from Neural Information Processing Systems (NIPS), December 5-8, Vancouver, British Columbia, Canada, 2005*, pp. 507–514.
- [17] J. Weston, S. Mukherjee, O. Chapelle, M. Pontil, T. A. Poggio, and V. Vapnik, "Feature selection for SVMs," in *Advances in Neural Information Processing Systems 13, Papers from Neural Information Processing Systems (NIPS), Denver, CO, USA, 2000*, pp. 668–674.
- [18] Y. Mohsenzadeh, H. Sheikhzadeh, A. M. Reza, N. Bathaee, and M. M. Kalayeh, "The relevance sample-feature machine: A sparse Bayesian learning approach to joint feature-sample selection," *IEEE Transactions on Cybernetics*, vol. 43, no. 6, pp. 2241 – 2254, 2013.
- [19] D. J. MacKay, "Bayesian interpolation," *Neural computation*, vol. 4, no. 3, pp. 415–447, 1992.
- [20] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York: Springer, 2006.
- [21] D. J. Newman, S. Hettich, C. L. Blake, and C. J. Merz, "UCI repository of machine learning databases," 1998. [Online]. Available: <http://www.ics.uci.edu/~mllearn/MLRepository.html>
- [22] G. Rätsch, T. Onoda, and K.-R. Müller, "Soft margins for AdaBoost," *Machine Learning*, vol. 42, no. 3, pp. 287–320, 2001.
- [23] C. J. Burges, "A tutorial on support vector machines for pattern recognition," *Data Mining and Knowledge Discovery*, vol. 2, no. 2, pp. 121–167, 1998.
- [24] G. C. Cawley, N. L. C. Talbot, and M. Girolami, "Sparse multinomial logistic regression via bayesian L1 regularisation," in *Proceedings of the Twentieth Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 4-7, 2006*, 2006, pp. 209–216.
- [25] R. Caruana and A. Niculescu-Mizil, "Data mining in metric space: an empirical analysis of supervised learning performance criteria," in *Proceedings of the tenth ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD)*. ACM, 2004, pp. 69–78.
- [26] U. Alon, N. Barkai, D. A. Notterman, K. Gish, S. Ybarra, D. Mack, and A. J. Levine, "Broad patterns of gene expression revealed by clustering analysis of tumor and normal colon tissues probed by oligonucleotide arrays," *Proceedings of the National Academy of Sciences*, vol. 96, no. 12, pp. 6745–6750, 1999.
- [27] T. R. Golub, D. K. Slonim, P. Tamayo, C. Huard, M. Gaasenbeek, J. P. Mesirov, H. Coller, M. L. Loh, J. R. Downing, M. A. Caligiuri, C. Bloomfield, and L. E.S., "Molecular classification of cancer: class discovery and class prediction by gene expression monitoring," *Science*, vol. 286, no. 5439, pp. 531–537, 1999.
- [28] B. Krishnapuram, L. Carin, and A. J. Hartemink, "Joint classifier and feature optimization for cancer diagnosis using gene expression data," in *Proceedings of the Seventh Annual International Conference on Research in Computational Molecular Biology*. ACM, 2003, pp. 167–175.
- [29] M. West, C. Blanchette, H. Dressman, E. Huang, S. Ishida, R. Spang, H. Zuzan, J. A. Olson, J. R. Marks, and J. R. Nevins, "Predicting the clinical status of human breast cancer by using gene expression profiles," *Proceedings of the National Academy of Sciences*, vol. 98, no. 20, pp. 11 462–11 467, 2001.
- [30] S. K. Shevade and S. S. Keerthi, "A simple and efficient algorithm for gene selection using sparse logistic regression," *Bioinformatics*, vol. 19, no. 17, pp. 2246–2253, 2003.
- [31] Y. Li, C. Campbell, and M. Tipping, "Bayesian automatic relevance determination algorithms for classifying gene expression data," *Bioinformatics*, vol. 18, no. 10, pp. 1332–1339, 2002.
- [32] R. Meir and T. Zhang, "Generalization error bounds for Bayesian mixture algorithms," *The Journal of Machine Learning Research*, vol. 4, pp. 839–860, 2003.

- [33] R. A. Choudrey, "Variational methods for Bayesian independent component analysis," Ph.D. dissertation, University of Oxford, 2002.
- [34] M. E. Tipping and A. Faul, "Analysis of sparse Bayesian learning," in *Proceedings of the Conference on Neural Information Processing Systems (NIPS)*, 2002, pp. 383–390.
- [35] P. Shivaswamy and T. Joachims, "Coactive learning," *Journal of Artificial Intelligence Research*, vol. 53, pp. 1–40, 2015.
- [36] B. Krishnapuram, L. Carin, M. A. Figueiredo, and A. J. Hartemink, "Sparse multinomial logistic regression: Fast algorithms and generalization bounds," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 6, pp. 957–968, 2005.
- [37] I. Psorakis, T. Damoulas, and M. Girolami, "Multiclass relevance vector machines: sparsity and accuracy," *IEEE Transactions on Neural Networks*, vol. 21, no. 10, pp. 1588–1598, 2010.
- [38] D. J. Wilkinson, "Bayesian methods in bioinformatics and computational systems biology," *Briefings in Bioinformatics*, vol. 8, no. 2, pp. 109–116, 2007.

APPENDIX

In order to compute a complete Bayesian classifier, feature and classifier co-learning includes computing this formula:

$$(\alpha, \beta) = \arg \max_{(\alpha, \beta)} p(\mathbf{t} | \alpha, \beta, \mathbf{S}) p(\alpha) p(\beta), \quad (18)$$

where we assume α and β are mutually independent. Equation 18 could be iteratively maximized between α and β . The re-estimation rules of α have been derived by [3]. In this appendix, we focus on deriving the re-estimating rules for β , which means that we need to compute the following equation:

$$\beta = \arg \max_{\beta} p(\mathbf{t} | \alpha^{\text{old}}, \beta, \mathbf{S}) p(\beta). \quad (19)$$

The hyperprior $p(\beta)$ follows the Gamma distribution, $p(\beta) = \prod_{k=1}^M \text{Gam}(\beta_k | c, d)$, where c and d are the parameters of the Gamma distribution.

Discussed in Section III-B, we can hardly calculate a closed form of marginal likelihood. Using Bayesian rules, the marginal likelihood is expanded as follows:

$$p(\mathbf{t} | \alpha^{\text{old}}, \beta, \mathbf{S}) = \frac{p(\mathbf{t} | \mathbf{w}, \theta, \mathbf{S}) p(\mathbf{w} | \alpha^{\text{old}}) p(\theta | \beta)}{p(\mathbf{w}, \theta | \mathbf{t}, \alpha^{\text{old}}, \beta)}. \quad (20)$$

Applying approximate Gaussian distributions for the sample and feature posteriors, in Section III-A, we can obtain $p(\mathbf{w}, \theta | \mathbf{t}, \alpha, \beta) \approx N(\mathbf{u}_\theta, \Sigma_\theta) * N(\mathbf{u}_w, \Sigma_w)$. As a result, we maximize the logarithm of (19):

$$\begin{aligned} L &= \log[p(\mathbf{t} | \alpha^{\text{old}}, \beta) p(\beta)] \\ &= \log p(\theta | \beta) - \log N(\mathbf{u}_\theta, \Sigma_\theta) + \log p(\beta) + \text{const} \\ &= \frac{1}{2} (\epsilon^T \mathbf{B}^{-1} \epsilon + \log |\mathbf{B}| - \log |\mathbf{H} + \mathbf{B}|) \\ &\quad + \sum_{k=1}^M (c \log \beta_k - d \beta_k) + \text{const}, \end{aligned} \quad (21)$$

where const is independent of β , $\epsilon = (\mathbf{D}^T(\mathbf{t} - \sigma) + \mathbf{k}_\theta)$ is an M -dimensional vector and $\mathbf{H} = \mathbf{D}^T \mathbf{C} \mathbf{D} + \mathbf{O}_\theta - \mathbf{E}$ is an $M \times M$ matrix. Practically, the latter two terms will disappear if we set $c = d = 0$.

To compute the optimal β , we first calculate differentiation of Equation (21):

$$\frac{\partial L}{\partial \beta_k} = -\frac{1}{2} \left(\frac{\epsilon_k^2}{\beta_k^2} - \frac{1}{\beta_k} + \frac{1}{\beta_k + h_k} - 2 \frac{c}{\beta_k} + 2d \right), \quad (22)$$

where h_k denotes the k th diagonal elements of \mathbf{H} . Note that $u_{\theta, k}^2 = \frac{\epsilon_k^2}{\beta_k^2}$ and $\Sigma_{\theta, kk} = \frac{1}{\beta_k + h_k}$, shown in (10) and (11). So, set (22) equal to 0. We obtain the updating formula for β :

$$\beta_k^{\text{new}} = \frac{2c + 1}{u_{\theta, k}^2 + \Sigma_{\theta, kk} + 2d}, \quad (23)$$

which is the same formula as the EM-based solution established by [2], and guarantees a local optimum. However, if using the methodology of Bayesian Occam's razor reported by MacKay in [19], we derive more efficient updating rules as follows:

$$\beta_k^{\text{new}} = \frac{\gamma_k + 2c}{u_{\theta, k}^2 + 2d}, \quad (24)$$

where $\gamma_k \equiv 1 - \beta_k \Sigma_{\theta, kk}$.